

REMARKS

Claims 1-8 and 11-20 are pending as of the Advisory Action mailed September 18, 2008. In that Advisory Action, the examiner entered the amendments requested by the applicant in the May 28, 2008 amendment after final action. That amendment had been filed after a representative of the applicant held a telephonic interview with the examiner. As recounted in the interview summary in that May 28, 2008 amendment, applicant's representative understood that an agreement was reached with the examiner with respect to one of the dependent claims, claim 10 – this claim was not obvious over Girson et al. in combination with Choi et al. Applicant correspondingly amended claim 1 to include the subject matter of claim 10 and believed that such a response was sufficient to overcome the final rejection. Upon receipt of the amendment after final, however, instead of removing the rejection of claim 10, as applicant understood would be the case, the examiner issued the September 18, 2008 Advisory Action rejecting the newly-amended claim 1. The case subsequently, and unintentionally, went abandoned as a result. A petition to revival has been filed contemporaneously herewith and should be granted.

Claims 1-4, 6-8, 11-17, and 19-20 stand rejected as anticipated by Girson et al. (U.S. Patent No. 7,111,179). The remaining claims each stand rejected under a proposed combination of Girson et al. and Choi et al. (U.S. Patent No. 6,233,690). Claim 9 was previously canceled. Claim 3 has been canceled above. Claims 22-26 have been added. No new matter has been added.

Applicant has amended both independent claims 1 and 13 above. The amendments clarify the reference to the two rate-based metrics being utilized to determine power policy (frequency and voltage regulation) on the machine, those metrics being instructions-per-clock cycle and memory references-per-cycle. In particular, the amendments recite that the runtime performance data includes instruction counts, memory references and cycle counts obtained from a timer interrupt in the performance monitor, which are then used to determine the two rate-based metrics. As amended, claim 1 recites:

1. (Currently Amended) An article comprising a machine-accessible medium having stored thereon instructions that, when executed by a machine, cause the machine to:

obtain from a performance monitor runtime performance data of a thread level utilization, wherein the runtime performance data is indicative of a set of execution characteristics of the thread including an instructions-per-clock cycle metric and a memory references-per-cycle metric, wherein the runtime performance data includes instruction counts, memory references and cycle counts; and

based on the runtime performance data, adjust an operating voltage or an operating frequency of the machine, wherein the operating voltage and operating frequency are nonzero.

Whether taken alone or in combination, neither Girson et al. nor Choi et al. teach or suggest the recited subject matter. In particular neither reference teaches, singularly or together, adjusting machine power based on a memory-references-per-cycle metric. Further, however, neither reference uses two rate-based metrics together to determine power scaling policy on a machine, as is claimed. Further still, neither reference teaches, singularly or together, obtaining runtime performance data that indicates instruction counts, memory references, and cycle counts, as claimed.

The office action points only to Choi et al., col. 3, lines 11-21 as purportedly teaching a memory references-per-cycle metric. Applicant respectfully, but strongly disagrees. Choi et al. is a single event triggering system, meaning that it looks to cycle down the processor clock if a particular triggering event results in a stall condition. Choi et al. decouples the system clock from an execution unit when a triggering event results in a long latency memory stall over a certain interval of time. Abstract. This decoupling allows the Choi et al. system to then power down the logic processor. Col. 2, l. 62 - Col. 3, 3. "The particular long latency stalls that are targeted for power reduction are those for which termination of the stall condition can be anticipated readily." Col. 3, l. 3-5. That is, Choi et al. decouples the clock and powers down the processor when the stall that is occurring is of a nature that the system will be able to determine, and predict even, when that stall will end. This way, the system will be able to quickly recouple the clock and resume powered up operation after the stall, but at least conserve power during the stall. Pointing to the language highlighted by the examiner, Choi et al. summarizes this point with the following example:

"When pending load stall is detected, power down is initiated following a selected interval that filters out shorter latency stalls. In this embodiment, a data return signal precedes the earliest release of the stall condition by a known interval." Col. 3, ll. 18-22.

Choi et al. nowhere mentions memory references-per-cycle. Choi et al. does not look to how many times a memory reference has been made, how many different memory references are made in a given period of time, whether different threads are simultaneously executing and making memory references, nothing of the sort. What the examiner cites in Choi et al. refers to a single trigger (e.g., an “a pending load operation” from a pending request) and whether that single operation has resulted in a stall in the amount of time the processor takes to respond. The stall that results from this single trigger may be used for power cycling, but that data does not indicate the actual memory references-per-clock cycle in the system, e.g., the number of times a memory reference has been made and by whom. Choi et al. is designed to take advantage of stall conditions, where close to no data transfer/processing is occurring, whereas the present application allows much more dynamic scaling when the machine is processing data, by allowing for power scaling based on metrics indicating that amount of desired data to process, namely instructions-per-clock cycle and memory references-per-clock cycle. These techniques, for example, allow a system to scale power up when parallel processes start running, i.e., when multiple threads resulting in many more memory references per cycle are executing simultaneously. Nothing in Choi et al. or any other art of record recognizes such benefits.

In any event, Choi et al. nowhere teaches the suggested subject matter, in particular using memory references-per-cycle to cycle power. Nor does Choi et al. teach combining such a metric with another rate-base metric, i.e., instructions-per-cycle, to cycle power. Of course, neither does Girson et al. In fact, none of the prior art provides a system able to achieve the features of the present application.

For at least the foregoing reasons, the rejections of claim 1 are traversed. Claim 1 is in condition for allowance.

The claims depending from claim 1 are in condition allowance by implication. However, applicant respectfully points out that various of these claims recite features that are worth highlighting as well. Claim 4, for example, has been amended to depend from claim 5, and further recite that a performance monitoring unit includes a plurality of counters that may simultaneously measure multiple different performance data metrics. The uniqueness and advantages of such a system are described in examples at the bottom of paragraph [0026] through paragraph [0028]. By having multiple dedicated counters, the system may

simultaneously monitor performance based events which in turn allows the system to observe “usage patterns” across application threads (not just a single stall from a single application) and simultaneously produce (if desired) multiple rate-based metrics that can be used for power scaling. None of the art of record describes counters used in the simultaneous way. None of the art of record is able to achieve advantages described in examples of the instant application.

Claim 13 is in condition for allowance for similar reasons to those outlined above with respect to claim 1, namely that claim 13 now recites:

13. (Currently Amended) A method comprising:

obtaining, from a performance monitor, runtime performance data of a thread level utilization, the runtime performance data being indicative of a set of execution characteristics of a thread, including an instructions-per-clock cycle metric and a memory references-per-clock cycle metric for a central processing unit (CPU) having an operating voltage and an operating frequency;

in response to the runtime performance data, determining if either the operating voltage or the operating frequency is at a desired value; and

in response to the determination, adjusting the operating voltage or the operating frequency in response to the instructions-per-clock cycle metric and the memory references-per-clock cycle metric, wherein the operating voltage and operating frequency are nonzero.

None of the prior art teaches or suggests obtaining a runtime performance data of a thread level utilization that indicates instructions-per-clock cycle and memory references-per-clock cycle for a CPU. Choi et al. is a single event triggering system, that decouples a clock during load stalls and other memory stalls, for the purposes of cycling down the power on a processor. There is no suggestion of triggering based on memory references. No counter is used to track memory references – the scoreboard (SB) 130 is a single bit flag that is set “when the associated register is awaiting a data return from memory” and thus clearly not a counter. Triggering is based on the length of a stall and whether the stall will end in a predictable timeframe – the number of memory references that occur during that stall is not tracked nor could one guess at the number of memory references that occur during a stall period. The number of clock cycles for the stall condition may be tracked, but not the number of memory references that result in a stall condition, only one load request triggers the stall condition and that load request presumably occurs in one cycle. Whether

multiple load requests, or multiple memory references, occur per cycle, is never measured or assessed.

The rejection of claim 13 and the claims depending are traversed.

Independent claim 22 was added by amendment above. The claim recites a method for adjusting operating voltage and/or operating frequency, notably by simultaneously monitoring multiple performance related events, such as instruction counts, memory references and cycle counts, which can then be used to determine rate-based performance metrics, such as instructions-per-clock cycle and memory references-per-clock cycle. As discussed in the present application, the current techniques (e.g., using a performance monitor as described) allow a machine to observe a number different events, simultaneously to determine runtime performance-patterns. For example, and as recited in claims depending from claim 22, systems can monitor different threads from the same application or threads from different applications altogether. Multiple performance counters in the PMU (performance monitor unit) may be used to achieve such tracking. See, e.g., paragraphs [0026] and [0027] of the present application. From here, appropriately power scaling policy may be determined. There is nothing in the prior art that teaches or suggests simultaneous monitoring, or the use of multiple calculated rate-based performance metrics to scale power.

For at least the foregoing reasons, claims 22-26 are in condition for allowance.

In light of the foregoing, applicant respectfully asserts that claims 1, 2, 4-8, 11-20, and 22-26 are in condition for immediate allowance. A prompt indication of allowability is earnestly solicited.

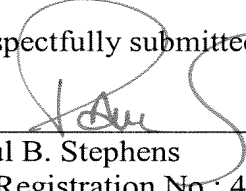
Should the examiner wish to discuss the foregoing, or any matter of form in an effort to advance this application toward allowance, he is urged to telephone the undersigned at the indicated number.

This paper is being submitted today along with a request for continued examiner and accompanying fee, and a petition for revival of an unintentional abandonment and accompanying fee. If it is determined that an additional fee is required, the Director is

hereby authorized to charge any deficiency in the fees filed, asserted to be filed or which should have been filed herewith to Deposit Account No. 13-2855, under Order No. 30320/17593.

Dated: May 22, 2009

Respectfully submitted,

By 
Paul B. Stephens

Registration No.: 47,970
MARSHALL, GERSTEIN & BORUN LLP
233 S. Wacker Drive, Suite 6300
Sears Tower
Chicago, Illinois 60606-6357
(312) 474-6300
Attorney for Applicant